

EVALUATING EXPLOITATION VERSUS EXPLORATION BY SIMULATION

John Debenham & Simeon Simoff
Faculty of IT
University of Technology, Sydney
PO Box 123 Broadway
Sydney, NSW, Australia
email: {debenham, simeon}@it.uts.edu.au

ABSTRACT

A general simulation model of market competition is developed to explore the effectiveness of and interactions between different types product exploration and exploitation strategies i.e. innovation, imitation and process improvement. The model, like real markets, is highly non-linear such that analytical solutions are not possible. We use simulation experiments to examine firm survival and the effectiveness of different strategy mixes and show how these depend on the length of time it takes for each strategy to bear fruit, the speed of new product diffusion and the duration of product life cycles. The model is freely available on the Internet and provides the basis for further experiments to examine the impact of different combinations of firm strategies on survival and performance.

KEY WORDS

Economic simulation. Evolutionary mechanisms.

1 Introduction

There are no simple analytical solutions to the investigation of exploitation and exploration because they involve highly non-linear dynamic systems for which analytical solutions do not exist, except in highly simplified and hence very unreal situations.

In order to examine alternative competitive regimes we distinguish between different types of competitors in terms of the extent to which they exploit current means of meeting customer needs and the extent to which they attempt to explore for new way of meeting customer needs (including new needs to meet). This distinction between exploitation and exploration lies at the heart of competitive strategy. Ansoff's (1965) classic competition matrix distinguishes between offering the same or different products to the same or different markets. Strategies focusing on the same products in the same markets are examples of exploitation strategies. All the rest involve some form of exploration, be that new means of serving existing markets, new markets for existing products or new product and market combinations.

The tradeoff between exploitation and exploration is a fundamental dimension of any strategy and is a concept which has been used to help understand the evolution of

biological species in ecological systems and their struggle for existence. For example, studies of social insects have revealed differences in the mix of exploration and exploitation among different types of ant colonies that have evolved in environments with more or less turbulence in the sources of food available [1]. In stable environments with relatively fixed sources of food more resources are devoted to the exploitation of known food sources and less to exploration whereas, in more dynamic environments, ant species evolve that devote more resources to exploration. In the same way firms in a market can adopt a mix of exploration vs exploitation strategies and the optimal balance or trade-off between the two will depend on the strategies of other competitors and the nature and dynamics of the market demand. Furthermore, algorithms based on social insect foraging behavior have been used to solve complex problems other types of approaches cannot solve [2].

James March [3] has summarized the key differences between exploitation and exploration strategies in the following way. "The essence of exploitation is the refinement and extension of existing competencies, technologies and paradigms. Its returns are positive, proximate and predictable. The essence of exploration is experimentation with new alternatives. Its returns are uncertain, distant and often negative". The strategies of exploitation versus exploration may be pursued in various ways.

The benefits of each strategy revolve around the probability of success of each type of strategy and costs and sacrifices involved. This in turn depends on the nature of the market and competitive situation, including the degree of turbulence and responsiveness of the environment and the strategies adopted by competitors. Determining an optimal strategy is thus complex because of uncertainties regarding the timing and payoffs of different strategies and because of non-linearities resulting from interactions among different firm strategies and of these strategies with market demand.

In order to examine the conditions under which different types of exploration and exploitation strategies are likely to succeed, we develop a model of competitive dynamics that allows firms to pursue different types of exploitation and exploration strategies in terms of devoting resources to improve the efficiency of supplying existing products (exploitation) as against using resources to develop new products (exploration). Two forms of explo-

ration are considered: (a) innovation in which a firm devotes resources to the invention of new products and (b) imitation in which the firm devotes resources to copying the products offered by other firms in the market. We use simulation techniques to examine the conditions under which different mixes of exploration, including both innovation and imitation, and exploitation perform best and how this compares to the situation of a firm that devotes no resources to exploitation or exploration and simply continues to supply the same types of products in the same way.

The simulation model developed must necessarily be a simplification of real world markets. But this is its purpose. By extracting from the real world key dynamic processes we are able to examine their role and impact in ways that are otherwise impossible. This approach to analysis is gaining increasing favour amongst scientists in many disciplines as they seek to understand the key processes underlying the development of economic, social and business systems [4], [5], [6], [7], [8]) and as researchers attempt to find solutions to complex nonlinear dynamic problems. As Chris Langton [6], one of the founders of the science of Artificial life observes in the context of research biology: “We trust implicitly that there are lawful regularities at work in the determination of this set [of realized entities], but it is unlikely that we will discover many of these regularities by restricting ourselves only to the set of biological entities that nature actually provided us with. Rather, such regularities will be found only by exploring the much larger set of possible biological entities”.

2 The model in outline

The model described here is designed to enable the performance of firms to be compared when they allocate resources different mixes of exploitation and exploration strategies. The exploitation strategies are (a) *producing* existing products with existing technology, or (b) *process improvement* i.e. using resources to improve the efficiency of production for existing products. The exploration strategies are (a) *innovation* i.e. using resources to discover new types of products, or (b) *imitation* i.e. using resources to copy new products produced by competitors.

The market model is of a “closed economy” — that is, a trading environment in which a fixed total amount of resources, modeled here in terms of a firm’s number of employees, is used by firms for generating product. These employees are also the source of market demand for the the product of firms. Using the money they receive for their labour services people purchase the product offered by firms in the market. All this takes place in successive, discrete time periods. At the beginning of each time period each firm has a budget for its labour. Each firm hires labour to the full extent of its labour budget. In the “final few moments” of each time period the following things happen:

- labour is paid by the firms in exchange for their work during that time period — at this stage labour has all

the money and the firms have none;

- the firms are paid by labour in exchange for the product — all product is either sold or written off before the next time period starts — at this stage the firms have all the money and labour has none;
- the firms are now “cashed up” and they commit all of their money by hiring labour for the next time period.

If no one buys a firm’s products in a period it receives no income. It will have spent all of its budget on hiring labour for that time period, will have nothing left for the next time period, and so it will go out of business. A firm’s profit in a time period is the amount that it receives for selling its product at the end of that time period less the amount that it spent on hiring labour at the beginning of that time period. If a firm makes a profit during a time period then its budget is increased in the next time period and so it will hire more labour than in the previous time period. If it makes a loss then its budget is decreased and size of its labour force contracts in the next time period. The objective of each firm is to survive. The total amount of money in the economy remains constant in time and is all placed on the table at the end of each time period as described above. The size of the labour force also remains constant as does the total and per capita remuneration that labour receives. At the beginning of each time period all money is committed by firms to hiring labour.

The firms differ in the way in which they allocate resources at the beginning of each time period to the four types of strategies. The four strategies are realised by allocating labour to four job types:

- workers who produce product — the proportion of firm i ’s money spent on workers is w_i .
- process improvers who improve work processes by generating “process knowledge” — that is knowledge of how to produce product better — the proportion of firm i ’s money spent on process improvers is p_i .
- imitators who design processes for producing products that have been discovered by other firms — the proportion of firm i ’s money spent on imitators is m_i .
- innovators who discover new products — the proportion of firm i ’s money spent on innovators is n_i .

If a firm discovers a new product during a time period then, at the end of that time period, other firms may decide to attempt to copy that product.

The objective of the simulation experiments described here is to understand the effect of values for the four basic variables w_i , p_i , m_i and n_i on a firm’s performance. These variables are constrained by:

$$0 \leq \{w_i, p_i, m_i, n_i\} \leq 1 \quad w_i + p_i + m_i + n_i = 1$$

for $i = 1, \dots, n$ where n is the number of firms.

3 The model in detail

The basic structure of the model, from the point of view of the economy, is shown in Figure 1. It owes much to [4]. At the beginning of each time period a labour force of fixed size is fully employed by a number of firms at a fixed wage rate. During each time period, the total costs for each firm are the amount it spends on hiring labour. The total costs for firm i are C_i . The total costs for all firms is $\sum_i C_i$, and this amount of money is entirely spent on hiring labour and so this is also the amount of money that the entire labour force will spend at the end of the time period when they purchase products. In each time period firm i allocates the effort of its workers across the range of products that firm i knows how to produce. That allocation of workers will lead — as determined by each product's process knowledge to the generation of actual product Q_i for firm i — where the underlining notation denotes a vector $Q_i = [q_{i,1}, q_{i,2}, q_{i,3}, \dots]$ — that is, $q_{i,j}$ is amount of the j 'th product that firm i produces in the time period. The total quantity of the j 'th product that is available at the end of the time period is $\sum_i q_{i,j} = q_j$. The total output, produced by all firms, at the end of the time period is represented as the vector $\underline{Q} = [q_1, q_2, q_3, \dots]$.

Innovation takes place when one firm begins producing a product that has not been produced before. For example, if: $\underline{Q} = [2, 3, 0, 2, | 0, 0, 0, 0, \dots]$ then this means that 2 units of product 1 are available, 3 units of product 2 and 2 units of product 4. New products discovered as a result of innovation are introduced to the right of the marker '|', and the marker is moved along so that it remains to the right of the most-recently-discovered product. So in the \underline{Q} shown above product number 3 is "out of production". Suppose that, in addition to the products in \underline{Q} , one of the firms is an innovator and that it commences production of a new product. This new product will be numbered 5.

Having reviewed the range of products that are available at the end of a time period, labour will have preferences over which particular products they desire. These preferences are expressed by labour attaching a relative demand measure [described below in Section 4] across the range of available products in \underline{Q} . The *relative demand* \underline{D} is used directly to determine the relative price per unit. For example, if the relative demand of product 1 is 0.8 and the relative demand of product 2 is 1.2 then the price per unit for product 2 will be 1.5 times the price per unit for product 1. At the end of each time period labour also places the total amount of money available, M , "on the table". M remains constant in time. The actual prices \underline{P} are set in proportion to the relative demand so as to clear the market. So the only way in which a product will be unsold is if its relative demand, is zero. For example, suppose that the total amount of money is 100, consider the following output and relative demand vectors: $\underline{Q} = [2, 3, 0, 2, | 0, 0, 0, \dots]$ $\underline{D} = [30, 30, 0, 0, 25, | 0, 0, 0, \dots]$ This will result in product 1 being sold at 15 per unit, product 2 being sold at 15 per unit and product 5 being sold at 12.5 per unit. The 2

Figure 1. The model from the point of view of the economy.


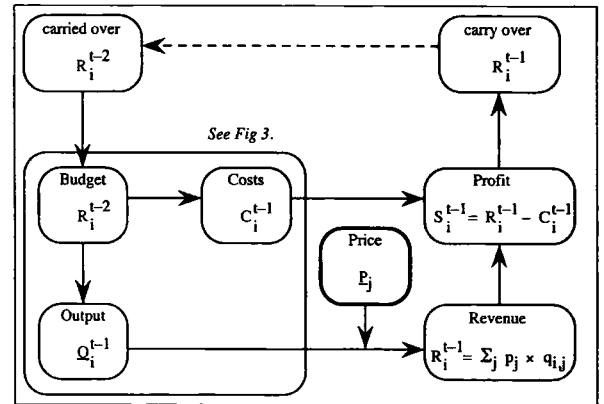


Figure 2. The model from the point of view of a particular firm i .



units of product 4 are unsold, and are written off by the firms that produced them. The total proceeds from selling at these prices is 100, which is also the total amount of money available.

Having determined the price vector \underline{P} , the model from the point of view of firm i is shown in Figure 2. Consider the time period $[t-1, t]$. At the beginning of this previous time period the firm will have carried over its revenue R_i^{t-2} derived in the previous time period and will have fully committed this revenue to hiring labour. The way in which the output vector Q_i^{t-1} and the costs C_i^{t-1} are determined for the products produced during the time period $[t-1, t]$ is described below in Figure 3. Having determined the output vector, and having calculated the price vector \underline{P}_i^{t-1} so as to clear the market as described above, the revenue for firm i , which is derived at the end of the time period $[t-1, t]$, is: $R_i^{t-1} = \sum_j (p_j \cdot q_{i,j})$. Hence the profit for this time period, S_i^{t-1} , is determined and so is the revenue that will be carried over to the next time period. The "anti-clockwise loop" shown in Figure 2 goes "round and round" from one time period to the next.

Figure 2 does not show how the carry over amount R_i^{t-2} , available at the start of time period $[t-1, t]$, generates output Q_i^{t-1} and costs C_i^{t-1} by the end of that time period. This is shown in Figure 3. The horizontal dashed

line in Figure 3 divides the figure into two time periods: $[t-2, t-1]$ in the upper part, and $[t-1, t]$ in the lower part. First, the carry over amount R_i^{t-2} from $[t-2, t-1]$ becomes the budget for the time period $[t-1, t]$. The budget R_i^{t-2} is entirely committed to hiring labour in the time period $[t-1, t]$. That is:

$$L_i^{t-1} = \frac{R_i^{t-2}}{c}$$

where c is the constant wage rate. For simplicity, c is set to unity. So a “unit of money” is the cost of a unit of labour for one time period. Labour is split in the proportions $w_i : p_i : m_i : n_i$ into the four categories workers, process improvers, imitators and innovators. The imitators attempt to build processes for producing products that have been discovered by other firms. If they are successful then they create a level of manufacturing expertise, or process knowledge, that is represented as a vector: Imi_i^{t-2} . For example: $\text{Imi}_i^{t-2} = [0, 0, 1.0, 0, 0, 0, 0, \dots]$ contains process knowledge with value 1.0 concerning product 3.

A firm's process improvers are allocated to improving the manufacturing processes for particular products. The i 'th firm's *process knowledge* is denoted by a vector \underline{A}_i . In the time period $[t-1, t]$ the process improvers may have found new process knowledge Pro_i^{t-1} — as for Imi_i^{t-2} this knowledge is represented as a vector denoting the product(s) that are the subject of the generated process knowledge. Likewise the innovators \underline{N}_i may discover process knowledge for new products, Inno_i^{t-1} . All knowledge generated during one time period may only be used in subsequent time periods, and so each firm's process knowledge available in the period $[t-1, t]$ is:

$$\underline{A}_i^{t-1} = \underline{A}_i^{t-2} + \text{Imi}_i^{t-2} + \text{Pro}_i^{t-2} + \text{Inno}_i^{t-2}$$

That is, each firm's process knowledge accumulates from one time to the next. It remains to describe how a firm's workers use this knowledge. Firm i 's workers are distributed across the range of products that the firm can produce as represented by the vector \underline{W}_i^{t-1} . The quantity of output that the workers generate in the time period is:

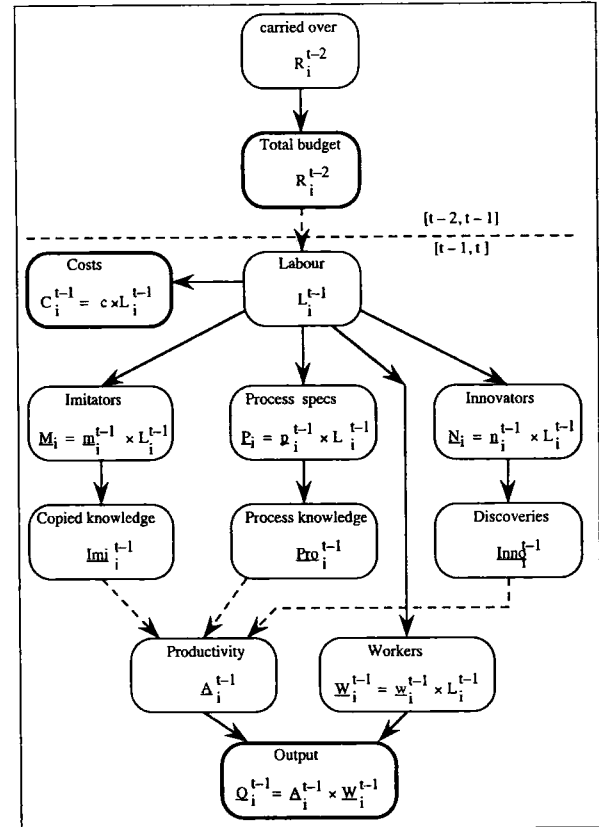
$$\underline{Q}_i^{t-1} = \underline{A}_i^{t-1} \cdot \underline{W}_i^{t-1}$$

where the \cdot symbol means that the vectors are multiplied together element by element.

4 Determining demand

The price of each type of product is determined at the end of each time period by the amount of product generated in that time period, by the total amount of money available, and by the “relative demand” for the different types of product which is determined by labour's preferences. Relative demand reflects the preferences of labour for different types of product. So a model of relative demand for each product is required to calculate unit price, as is a model of

Figure 3. An allocation of resources leads to output and costs for firm i . The dashed line separates two time periods, and dashed arrows mean that the new knowledge is not available until the following time period.



supply — ie: the product generated. Relative demand is considered now, and supply is considered in the next sub-section.

A modified Bass model [9] with repeat purchase is used to model relative demand for different products in the market subject to a fixed total overall market demand. The rate of new product diffusion the rate of repurchase depends on the type of product or service, as numerous studies of new product diffusion and adoption have indicated. Thus the rate of development of the demand for automobiles is not the same as that for a new beverage because at best each member of the population will purchase one or two automobiles but may purchase a beverage repeatedly. The type of products which we have in mind in developing our model are packaged food product in a market with fixed total demand. In each time period there is a total demand for a fixed σ units of product (eg: σ packaged dinners). σ is called the market size. Given a particular product (eg: a particular packaged dinner), in a particular time period $[t-1, t]$, the *initial penetration*, P^{t-1} , is the size of the population who has purchased this product at least once either during or before this time period. In time

period $[t-1, t]$, the *first-time sales*, N^{t-1} , are sales made of this product during this period to those who have not purchased this product previously. Suppose that the growth of initial penetration t is proportional, for some penetration constant g , to the size of the population that has yet to purchase this product. Then initial penetration in time period $[t-1, t]$, P^i , satisfies: $P^0 = \gamma \cdot m$, $P^1 - P^0 = \gamma \cdot (\mu P^0)$, $P^2 - P^1 = \gamma \cdot (\mu P^1)$. Or as a continuous approximation:

$$\frac{dP}{dt} = \gamma \cdot [\mu - P].$$

Solving this differential equation gives the initial penetration: $P^t = \mu \cdot (1 - \exp(-\gamma t))$. First-time sales is the rate of change of initial penetration. So if N^t is first time sales at time t : $N^t = P^t - P^{t-1}$, and as a continuous approximation:

$$N^t = \frac{dP^t}{dt} = \mu \cdot \gamma \cdot \exp(-\gamma t). \quad (1)$$

Now suppose that once labour has purchased a product, labour continues to purchase that product with a probability of α . That is, if T^i is total sales in time period $[i-1, i]$: $T^{i+1} = N^{i+1} + \alpha \cdot T^i$ where N^i is first time sales in time period $[i-1, i]$. Then: $T^0 = N^0$, $T^1 = \alpha \cdot N^0 + N^1$, $T^2 = \alpha^2 \cdot N^0 + \alpha \cdot N^1 + N^2$, etc. Or as a continuous approximation: $T^t = \int_{i=0}^t \alpha^{t-i} \cdot N^i \cdot di$. Evaluating this using equation 1:

$$T^t = \frac{\mu \cdot \gamma}{\ln(\alpha) + \gamma} [\alpha^t - \exp(t \cdot \gamma)]. \quad (2)$$

Which, for a market size of $\sigma = 100$ gives total sales values for each time period as shown in Figure 4 for various γ and α . The sales graphs in Figure 4 are now used to model relative demand. The discovery of a new product by an innovating firm can lead to a substantial shifts in demand for different firm's products, depending on how rapidly the new product diffuses through the market and the peak demand achieved, which depend on the values of the parameters a and g . For example, with $\gamma = 0.2$ and $\alpha = 0.9$, there is a rapid growth in demand for a new product to nearly 50% share within 8 time periods. The choice of α and γ in the simulations described below substantially effects the speed and extent of new product diffusion and the duration of the life cycle of the product, as can be seen from the Illustration in Figure 4. This affects the results of firm's using different strategies as we will show.

For a given market size σ , equation 2 has two variables: α and γ . Given the values of a total sales function in the first two time periods, f^0 and f^1 , it is easy to calculate α and γ :

$$\gamma = \frac{f^0}{\sigma}, \quad \alpha = \frac{f^1}{f^0} - \frac{\sigma - f^0}{\sigma}$$

and so knowing the first two values of a total sales function is to know "all there is" about it.

Returning now to the problem of modelling relative demand. The general shape of the total sales function in

Figure 4. Total sales for each time period for a market of size $\mu = 100$ and various γ and α . From top-left to bottom-right: $[\gamma = 0.1, \alpha = 0.7]$, $[\gamma = 0.1, \alpha = 0.9]$, $[\gamma = 0.2, \alpha = 0.7]$, $[\gamma = 0.2, \alpha = 0.9]$





Figure 4 is a fair description of how interest in a new product, such as packaged foodstuffs, might be expected to develop. Equation 2, for some values of α and γ is used here to model relative demand. So each product has a relative demand determined by equation 2, with its own values of g and a , for some fixed arbitrary σ , say, $\sigma = 1$. Consumers distribute their money over the different products in proportion to their relative demand D for each product as described above. *The prices per unit of the products is in proportion to their relative demand, and are set so as to clear the market.*

5 The Return on Investment of Different Strategies

The return on investment (ROI) of different strategies may be indicated in terms of the area beneath the relative demand curve for a product. For example an innovator invests for a period of time and discovers a new product. That directly benefits the innovator until an imitator learns to imitate that product, at which time the innovator will share the benefit with the imitator. The innovators relative benefit is shown as the hashed area in the left of Figure 5. Likewise the imitators relative benefit is shown in the middle of Figure 5.

If an innovating firm is competing with an imitating firm then, when the innovating firm discovers a new product, it benefits entirely from the revenue derived from selling that product until the imitating firm learns how to produce that product and thereafter the two firms share the revenue from that product. So the difference between the sales volumes derived by the innovating firm and the imitating firm is related to the area beneath the being-imitated product's relative demand curve from its beginning to the time

Figure 5. Return on investment for an innovator, an imitator and an improver.



at which the imitating firm learns how to imitate that product. The area beneath a relative demand curve (2) from the beginning to time t is:

$$\int_{x=0}^t \frac{\mu \cdot \gamma}{\ln(\alpha) + \gamma} \cdot (\alpha^x - \exp(-x \cdot \gamma)) \cdot dx = \frac{\mu \cdot \gamma}{\ln(\alpha) + \gamma} \cdot \left(\frac{\alpha^t}{\ln(\alpha)} + \frac{\exp(-t \cdot \gamma)}{\gamma} - \frac{1}{\ln(\alpha)} - \frac{1}{\gamma} \right) \quad (3)$$

which gives an area under the entire curve of:

$$\frac{\mu \cdot \gamma}{\ln(\alpha) + \gamma} \cdot \left(-\frac{1}{\ln(\alpha)} - \frac{1}{\gamma} \right)$$

For example if $\gamma = 0.1$ and $\alpha = 0.7$ then the area under the entire curve is 448.142, and the expression (3) tends asymptotically to this value.

The process improver will invest in improving production for a recently discovered, or copied, product. When the investment in process improvement exceeds the improvement threshold the process knowledge for that product will increase by 1.0. This is illustrated in the right of Figure 5.

What Figure 5 shows is that if an innovating firm, an imitating firm and an improving firm are coexisting in a moderately stable way then we expect the innovation threshold to be greater than the process improvement threshold, which in turn will be greater than the imitation threshold. This turns out to be the case.

6 Model Implementation on the Internet

The "economy" described above has been implemented as a Java applet and is available on the World Wide Web at: <http://www-staff.it.uts.edu.au/~debenham/research/evolution1/>

The use of Microsoft Internet Explorer with Java enabled is recommended. It has been used to conduct the experiments reported herein.

7 Conclusion

We have shown how the choice of exploration versus exploitation strategies is a complex problem without any analytical solutions. This arises in part because innovation,

imitation and process improvement are not deterministic processes and because of interactions among the strategies of different firms. This makes the whole market system a highly non-linear one. As a result we have to utilise simulation models to examine the conditions under which different strategies are successful or not in terms of firm and competitor survival and ROI. We have shown how survival depends on the timing of innovation, imitation and process improvement and the speed of diffusion and level of penetration of products. These factors affect the trade-off between the more shorter term gains from exploiting existing or recently developed new products against the more distant gains from developing new products. Optimal rates of allocating resources have been detected under different threshold and market demand conditions in markets with two firms competing and these appear to dominate any imitation strategies.

References

- [1] Bonabeau, E., Theraulaz, G., Deneubourg, J.: Dominance orders in animal societies — the self-organization hypothesis revisited. *Bulletin of Mathematical Biology* **61** (1999) 727 — 757
- [2] Bonabeau, E., Meyer, C.: Swarm intelligence: A whole new way to think about business. *Harvard Business Review* (2001) 107 — 114
- [3] March, J.: A Preface to understanding how decisions happen in organizations. In: *Organizational Decision-Making*. Cambridge University Press (1996)
- [4] Andersen, E.: Evolutionary econometrics: from Joseph Schumpeter's failed econometrics to George Price's general evometrics and beyond. In: *Proceedings Third Workshop in the Economic Transformation of Europe*. (2004)
- [5] Casti, J.: *Would-Be Worlds: How Simulation is Changing the Frontiers of Science*. John Wiley and Sons, Inc., New York (1997)
- [6] Langton, C.: *Artificial Life: An Overview*. The MIT Press, Cambridge, MA. (1997)
- [7] Iansiti, M., Levien, R.: *The Keystone Advantage: What the New Dynamics of Business Ecosystems Mean for Strategy, Innovation, and Sustainability*. Harvard Business School Press (2004)
- [8] Tesfatsion, L.: How Economists Can Get Alife. In: *The Economy as an Evolving Complex System II*. Addison-Wesley, Reading, MA (1997) 534 — 564
- [9] Shikhar, S., Mahajan, V.: The effect of reward structures on the performance of cross-functional product development teams. *Journal of Marketing* **65** (2001) 35 — 53